

# ***Scalable Visualization Clusters***

## *using commodity hardware*

Alan Heirich, Hewlett-Packard

[alan.heirich@hp.com](mailto:alan.heirich@hp.com)

ACTS workshop on the DOE Advanced  
Computational Software Collection

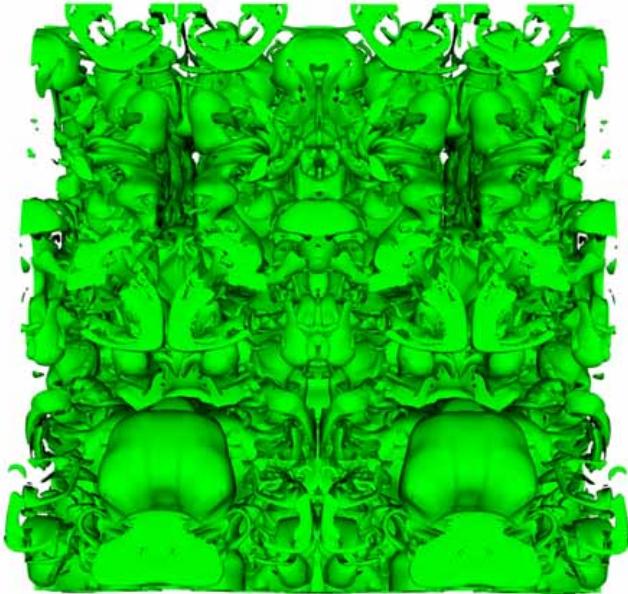
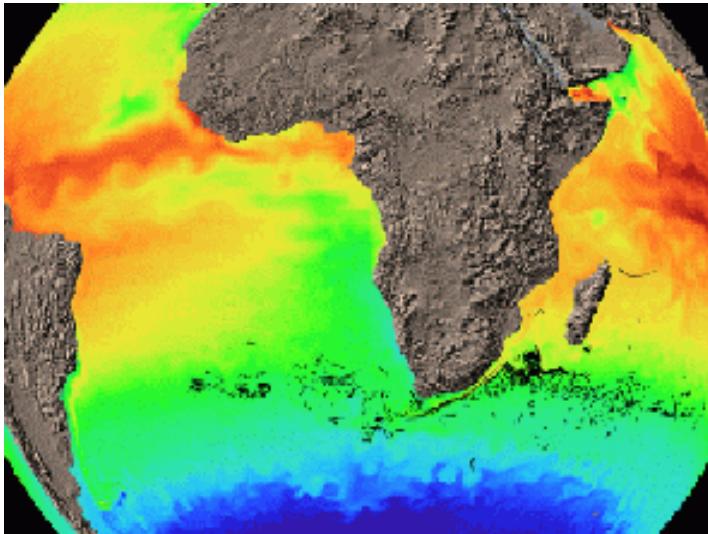
September 7, 2002

Lawrence Berkeley National Laboratory

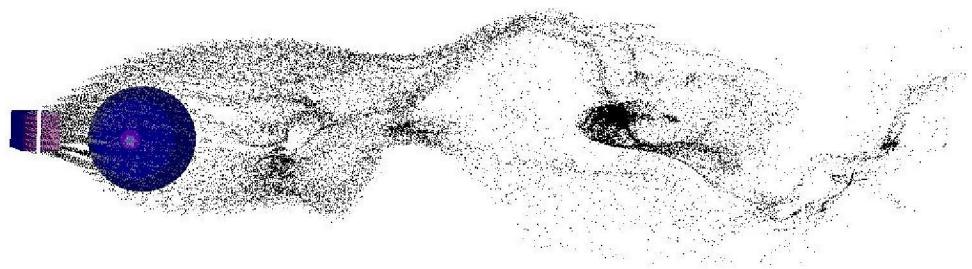


# References – in proceedings

- IEEE Parallel and Large-Data Visualization and Graphics Symposium (2001)  
*Scalable Interactive Volume Rendering Using Off-the-Shelf Components*  
S. Lombeyda, L. Moll, M. Shand, D. Breen and A. Heirich
- Eurographics 4<sup>th</sup> Workshop on Parallel Graphics and Visualization (2002)  
*Distributed Rendering of Interactive Soft Shadows*  
M. Isard, M. Shand and A. Heirich
- IEEE Parallel Visualization and Graphics Symposium (1999)  
*Scalable Distributed Visualization Using Off-the-Shelf Components*  
A. Heirich and L. Moll
- IEEE Symposium on Field Programmable Custom Computing Machines (1999)  
*Sepia: Scalable 3D Compositing Using PCI Pamette*  
L. Moll, A. Heirich, and M. Shand



*Los Alamos National Laboratories, Advanced Computing Laboratory*



*Bruckschen, Kuester, Hamann and Joy (2001)*

*“Real time out-of-core visualization of particle traces”*

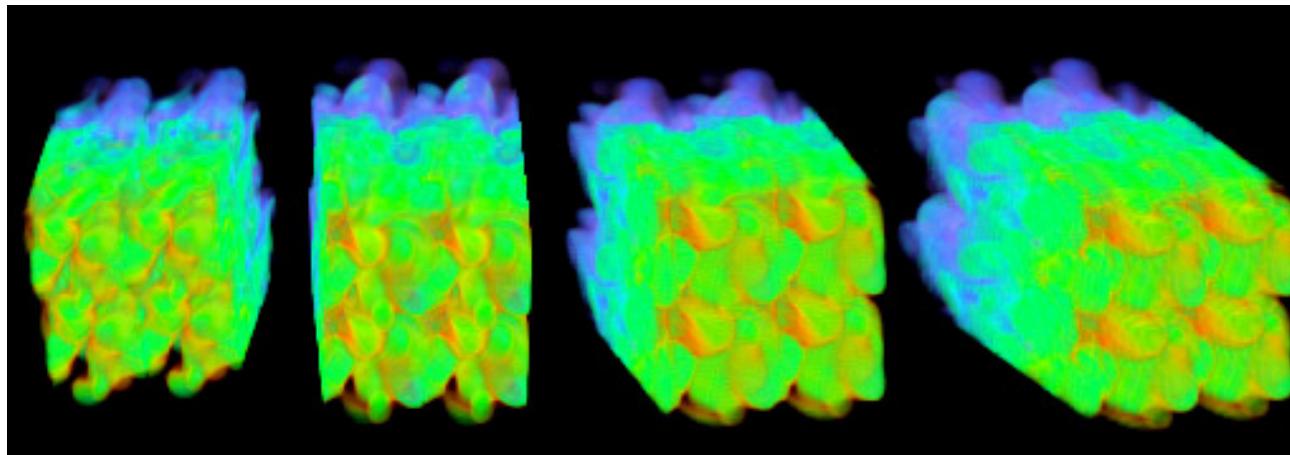
*surface rendering, reflection, commutative*



**Demonstration:**  
**512<sup>3</sup> voxels**  
**1024<sup>2</sup> pixels**  
**24-28 fps**  
**8 computers**



*volume rendering,  
transmission,  
non-commutative,  
(associative)*

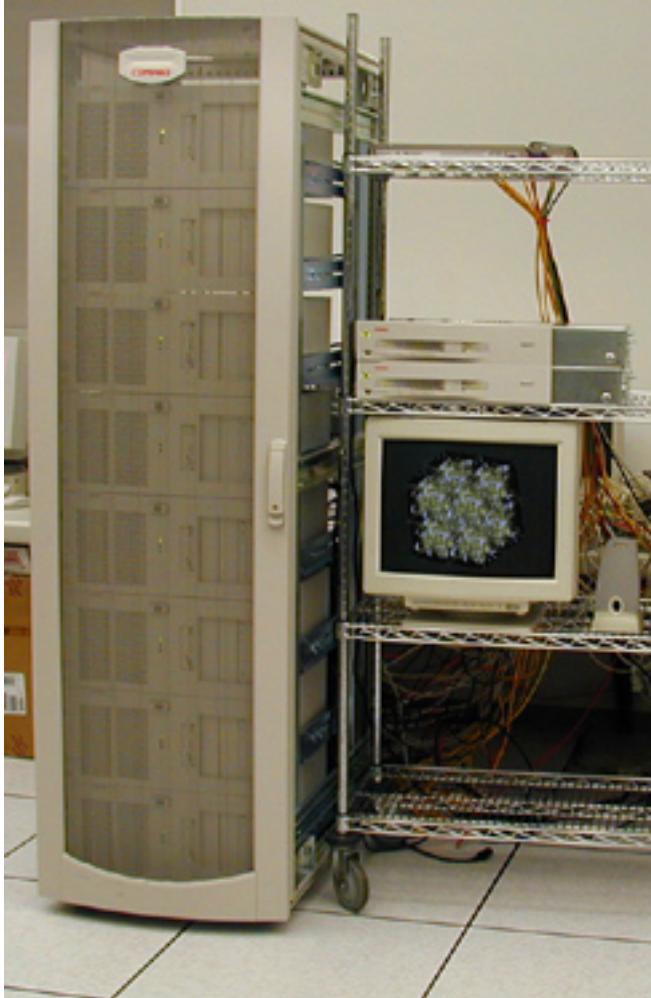


*Lombeyda, Shand, Moll, Breen & Heirich, PVG2001*

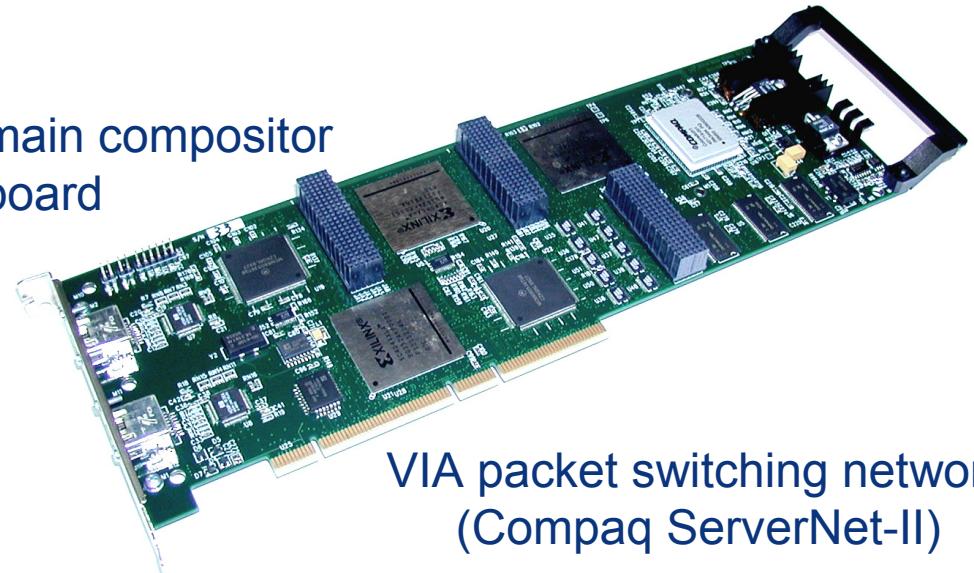


# Sepia-2a prototype (2002)

*compositing network with programmable arithmetic*

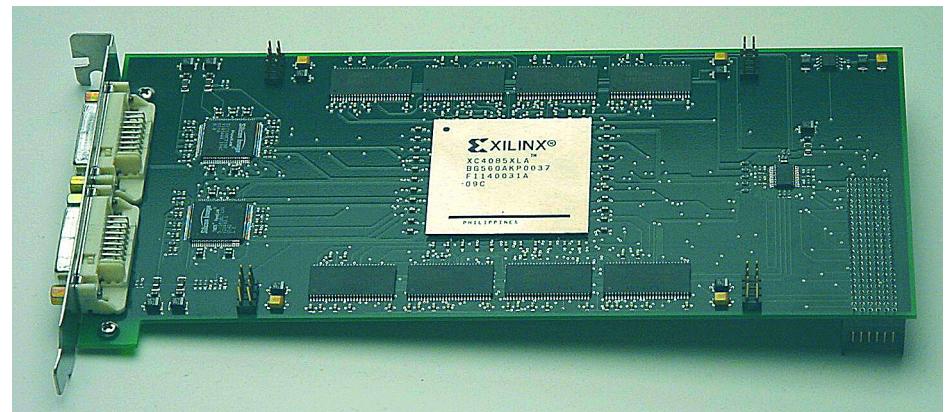


main compositor  
board



VIA packet switching network  
(Compaq ServerNet-II)

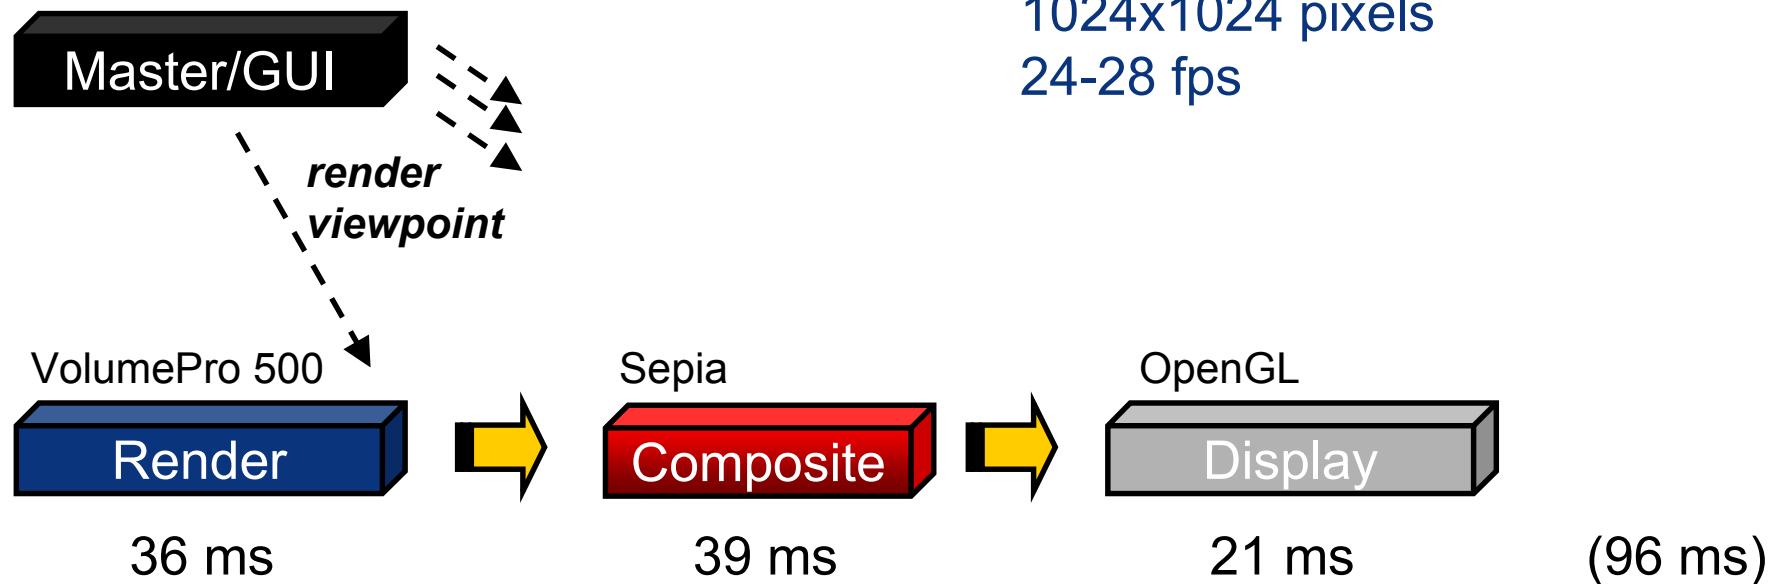
DVI I/O daughterboard



# Sepia ray-casting 2001

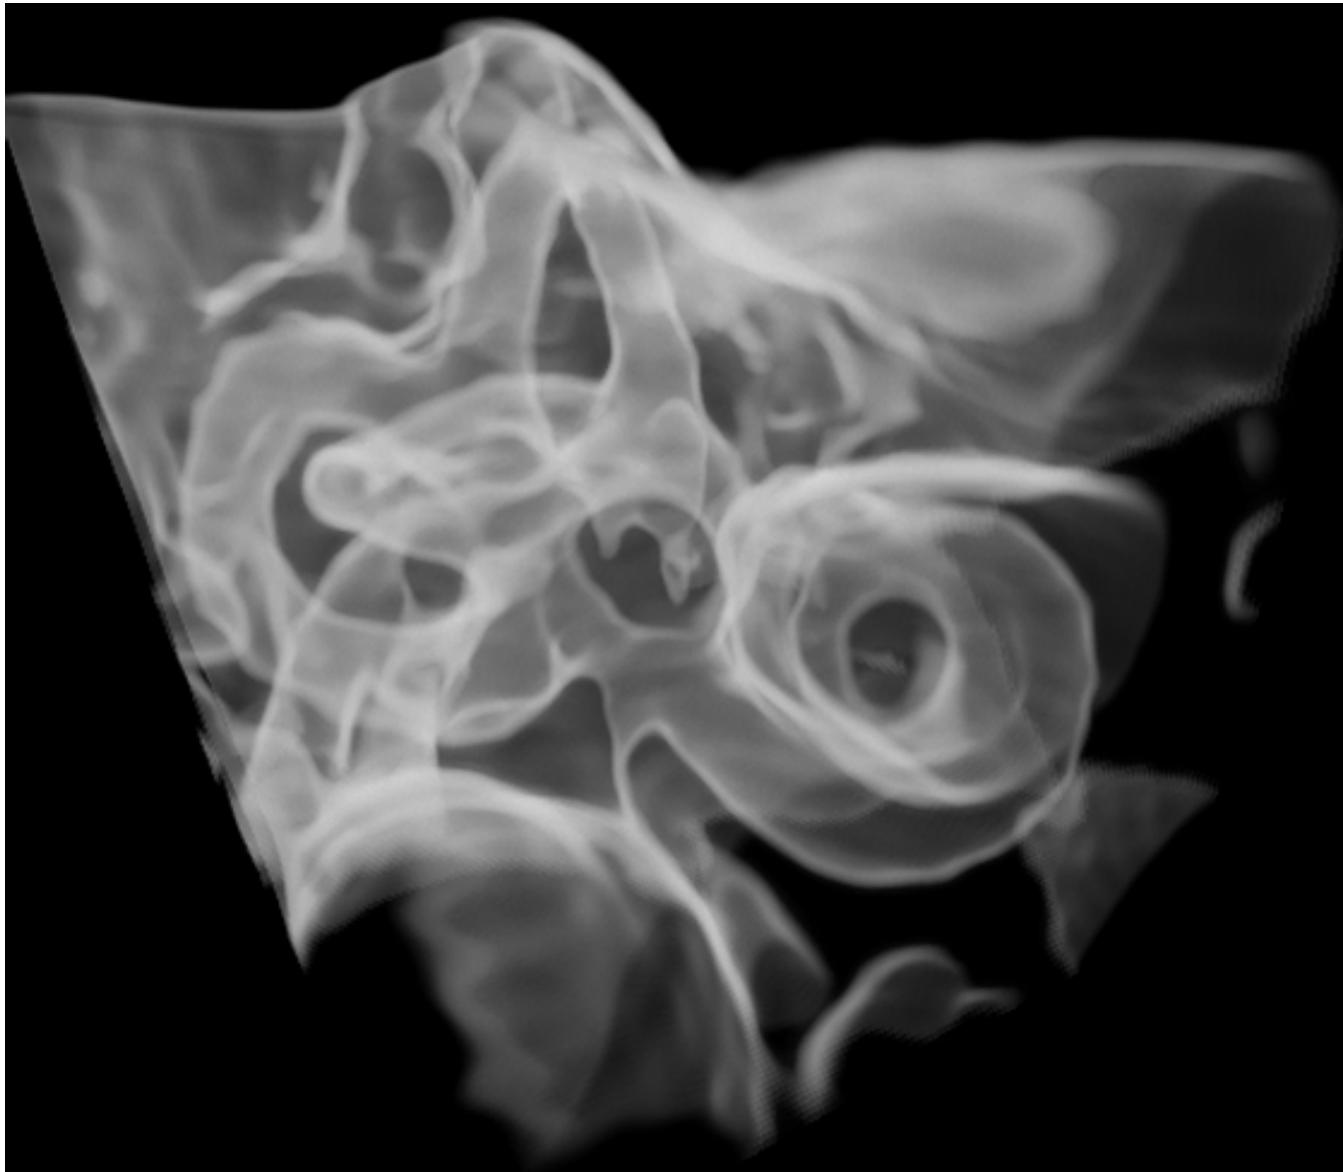
subvolume bricks 256x256x256  
500 MV/s ray-casting engine  
shear-warp algorithm  
master/slave application  
8+1 Pentium-III workstations

3.5 GV/s sustained  
512x512x512 voxels  
1024x1024 pixels  
24-28 fps



**Figure-of-Merit FOM = (Data MV) \* (Image MP) \* (Frames per second)**

Name	Yr	FOM (x 10 <sup>-12</sup> )	GV/ s	Platform	GPU	GPU #	Data MV	Image MP	FPS
<b>Binary swap <i>Ma-LANL</i></b>	94	<b>1</b>	.004	CM-5 Sparc	Software ray-cast CM-5	512	16	0.25	0.25
<b>Visual Super-computer <i>Ma-Mitsubis</i></b>	01	<b>380</b>	0.6	Pentium-III	TeraRec Volume Pro 500	8	128	0.5625	4.8
<b>Chromiu m <i>Stanford-LLNL</i></b>	02	<b>528</b>	1.87	Pentium-III	NVIDIA GeForce -3	16	128	0.25	15
<b>Sepia ray-cast <i>HP-Caltech</i></b>	01	<b>4,000</b>	3.5	Pentium-III	TeraRec Volume Pro 500	8	128	1.0	28
<b>Sepia OpenGL</b>	02	<b>tbd</b>	tbd	Pentium-IV	ATI or NVIDIA	8-12	128	1.831	tbd
<b>TRex <i>LANL</i></b>	01	<b>10,000</b>	5	Origin2K MIPS	SGI Infinite Reality-2	16	1024	1.831 (inferred)	5



*Engel, Kraus & Ertl (2001)*

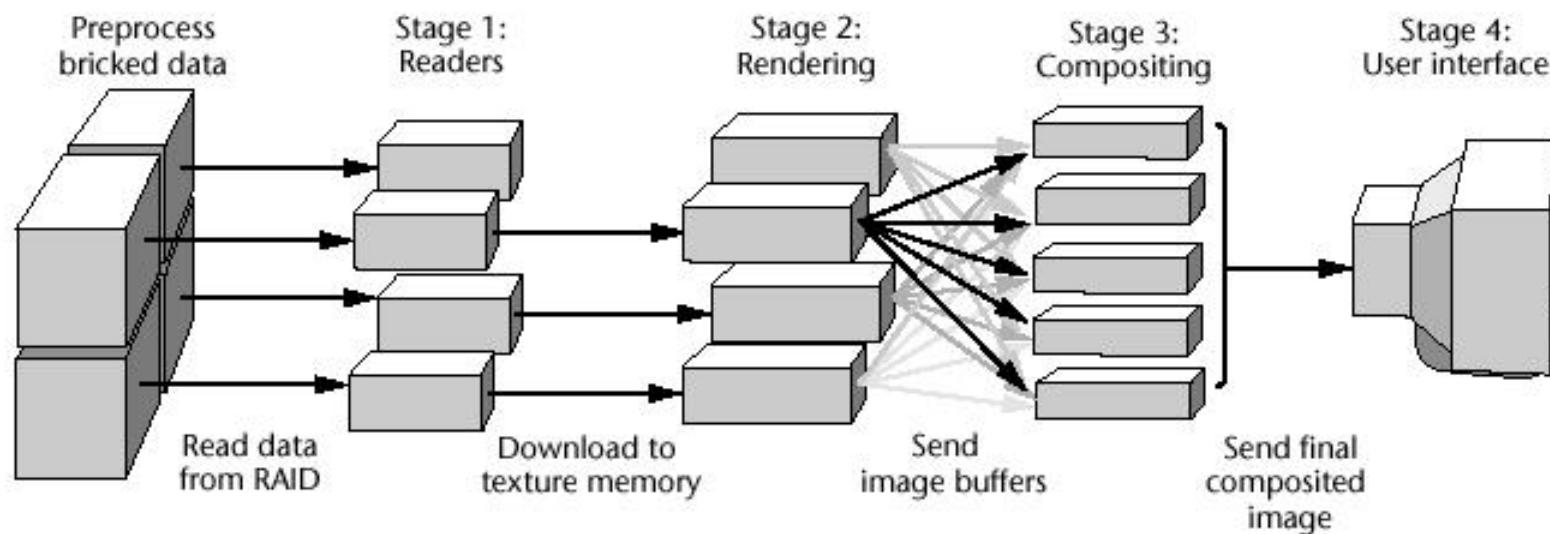
*"High quality pre-integrated volume rendering  
using hardware-accelerated pixel shading"*



# TREX texture mapping 2001

OpenGL 3D texture mapping  
three threads per accelerator  
SGI Origin2000, 128 CPUs  
16 InfiniteReality-2 OpenGL

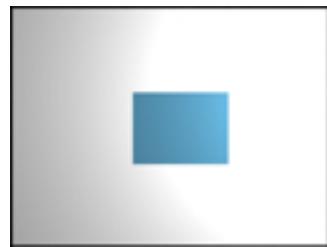
5 GV/s sustained  
1024x1024x1024 voxels  
5 fps



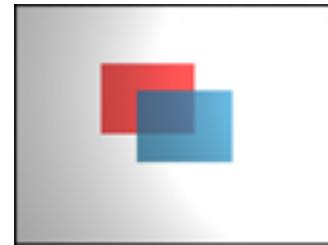
*Kniss, McCormick, McPherson, Ahrens, Painter, Keahey and Hansen, IEEE CG&A 2001*



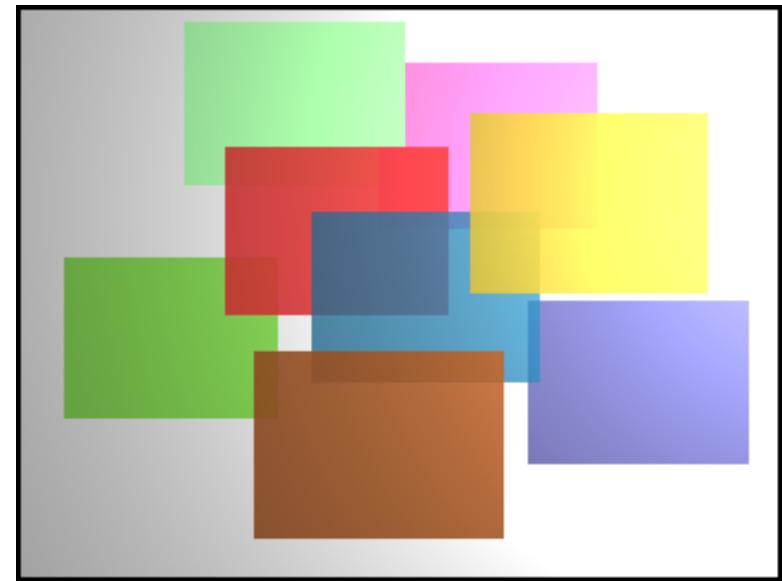
# Programming model – Virtual Frame Buffer



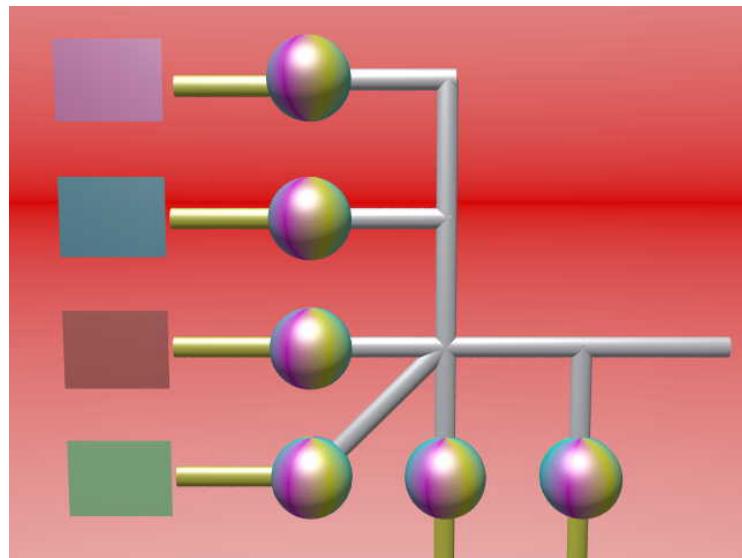
I/O = 1  
Composites = 0



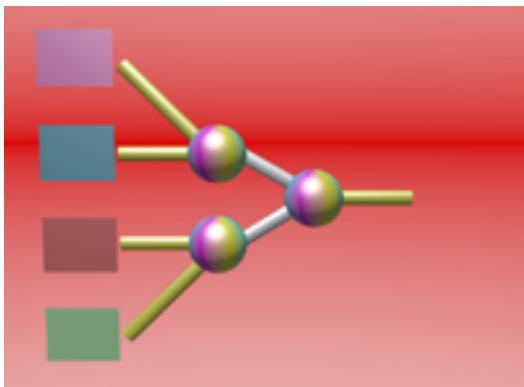
I/O = 2  
Composites < 1



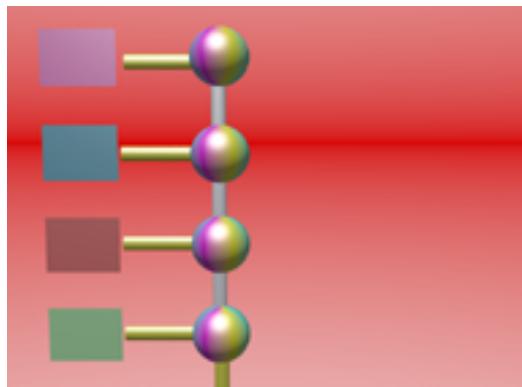
I/O = n  
Composites < n



**HP Sepia O(n)**



**Sony GScube O(n)**



**HP PixelFlow O(n)  
(UNC PixelFlow)**

**Sepia with Clos topology:**

n compositors

cn switches, p ports each

if p=16:

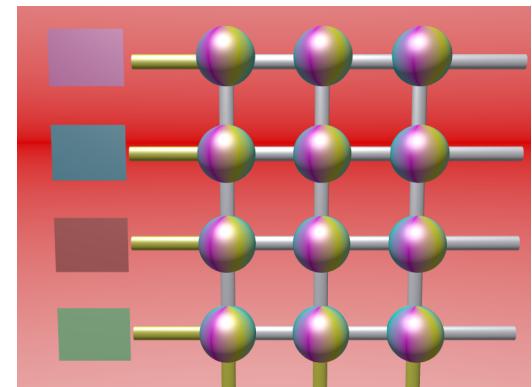
c = c1 for n in 1..128

c = c2 for n in 129..8192

**Sepia with linear topology:**

n compositors

n / (p-2) switches

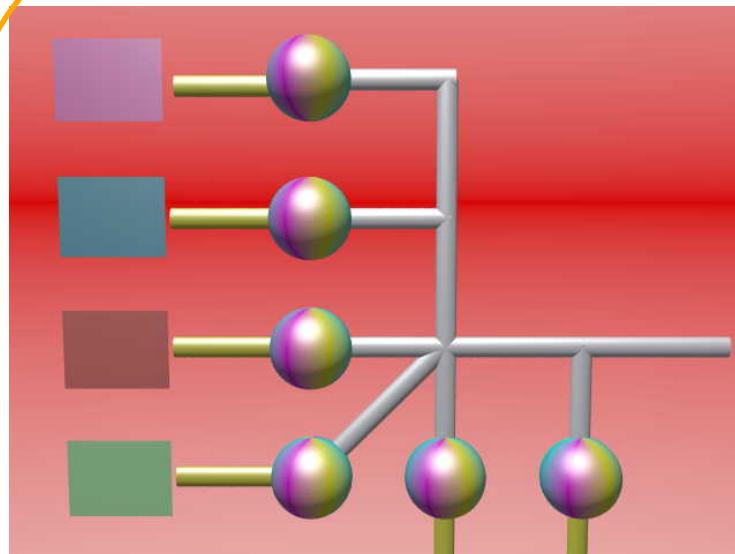


**Intel Illuminate O( $n^2$ )  
(Stanford Lightning-2)**



*Cluster  
architecture*

*Embedded  
architectures*



**HP Sepia O(n)**

**Sepia with Clos topology:**

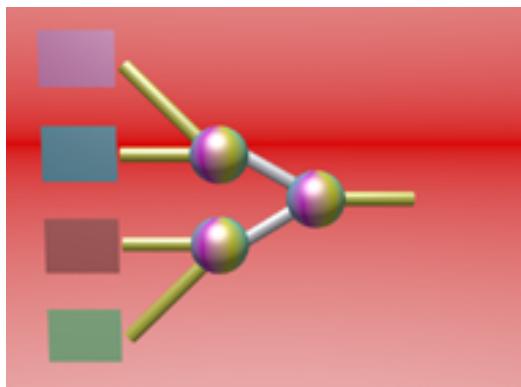
n compositors  
cn switches, p ports each

if p=16:

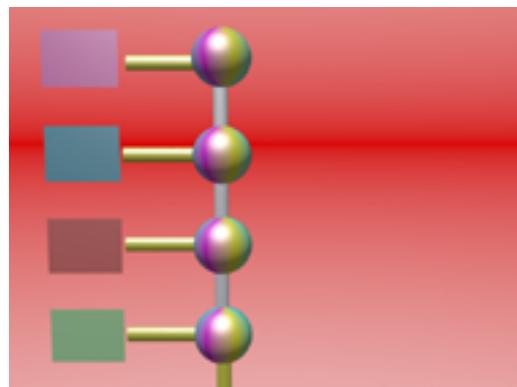
c = c1 for n in 1..128  
c = c2 for n in 129..8192

**Sepia with linear topology:**

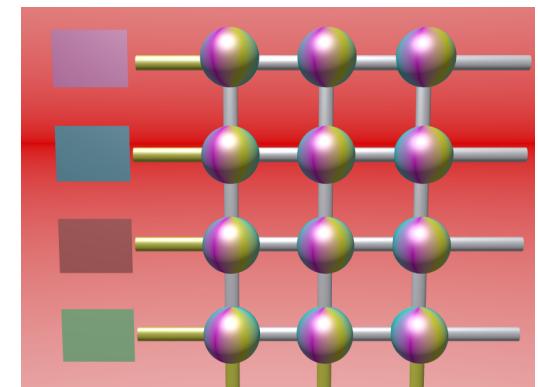
n compositors  
 $n / (p-2)$  switches



**Sony GScube O(n)**



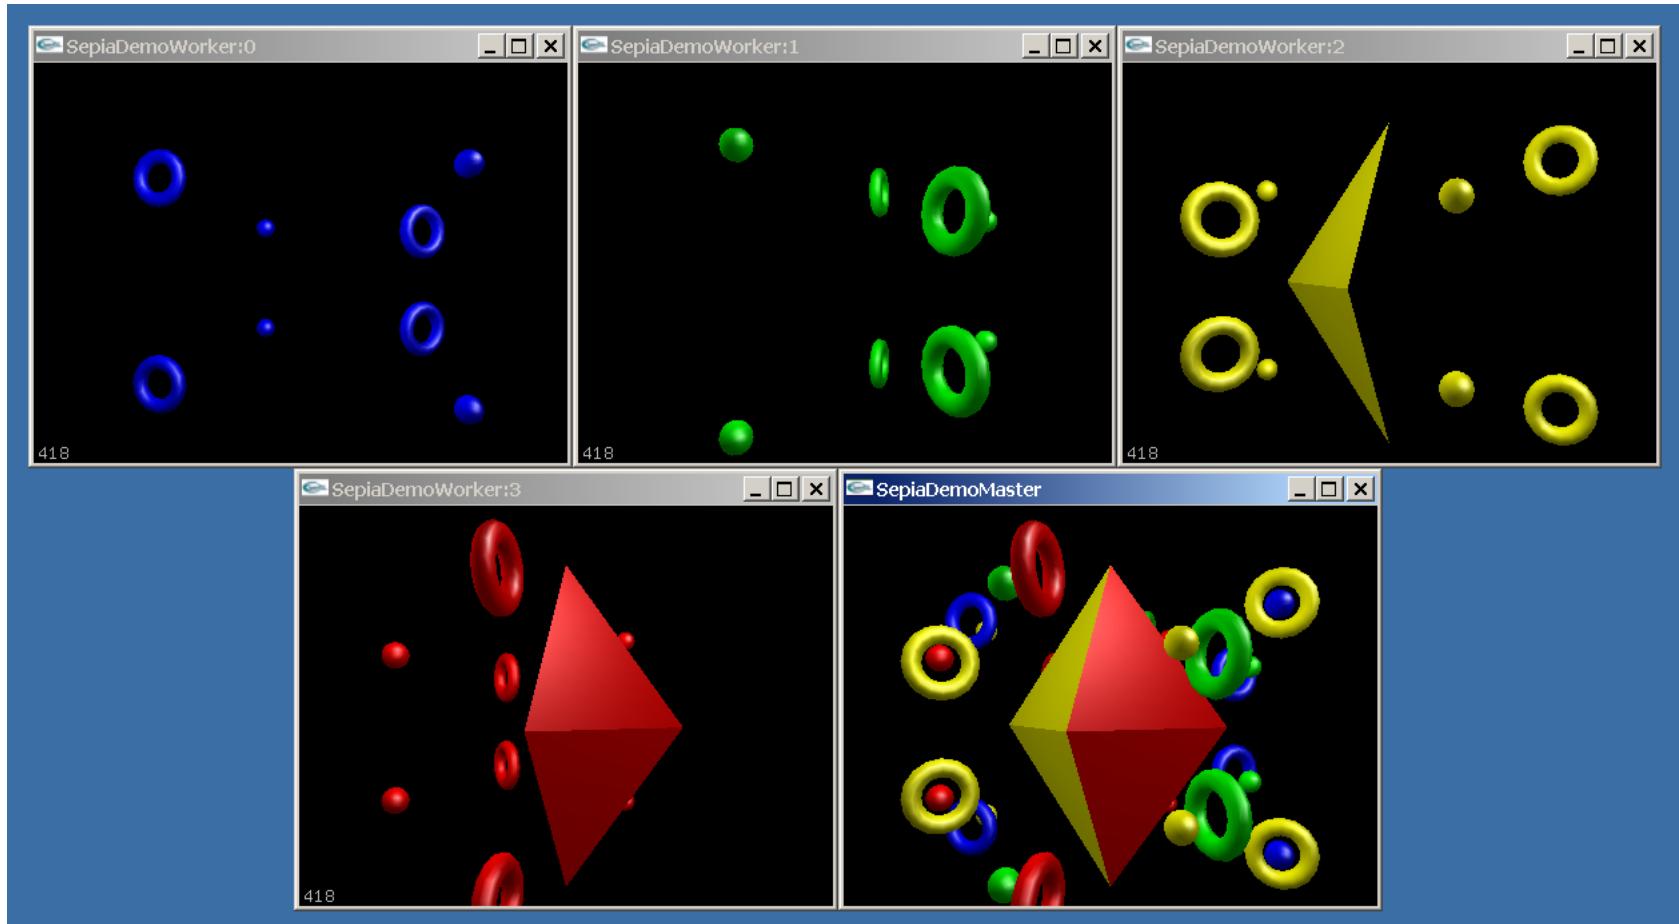
**HP PixelFlow O(n)  
(UNC PixelFlow)**



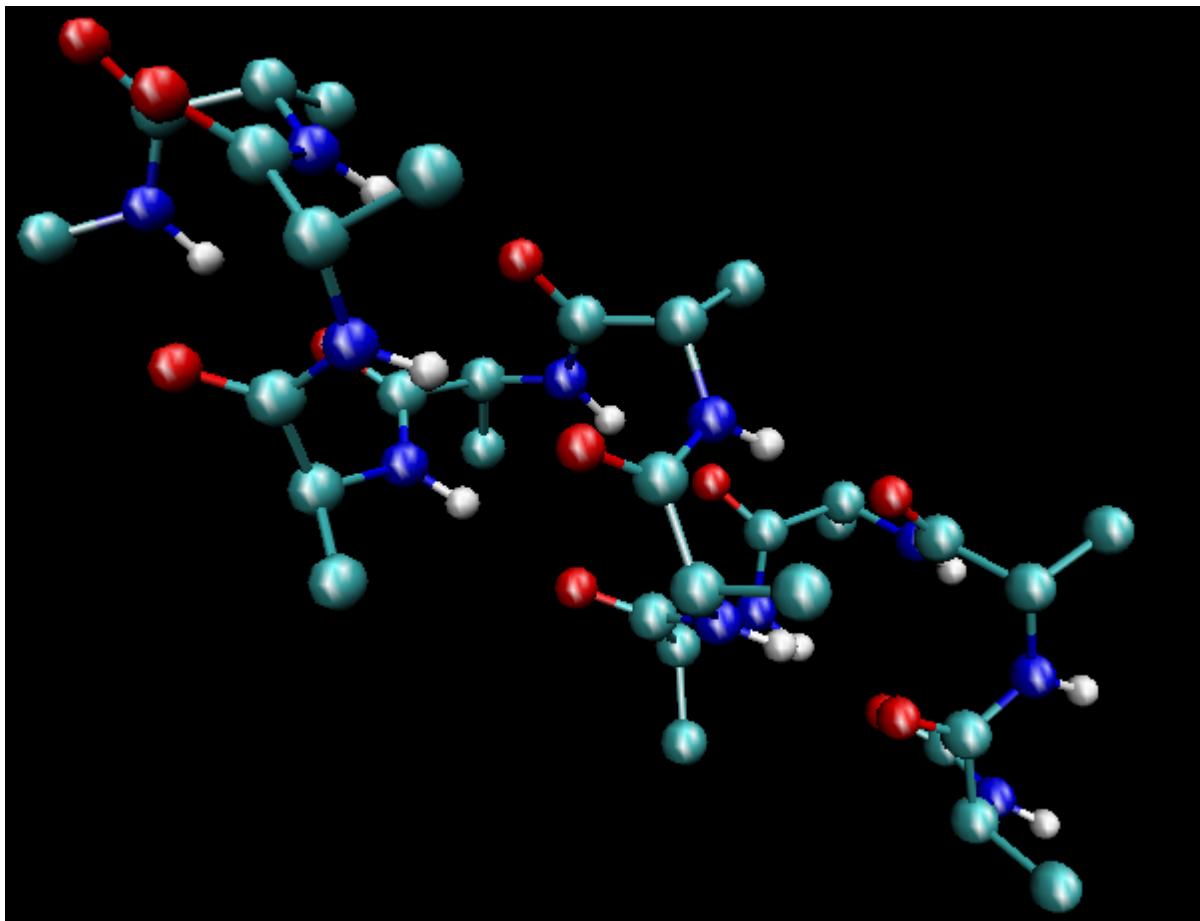
**Intel Illuminate O( $n^2$ )  
(Stanford Lightning-2)**

# Depth-compositing – illustration

Models light transmission (commutative arithmetic)

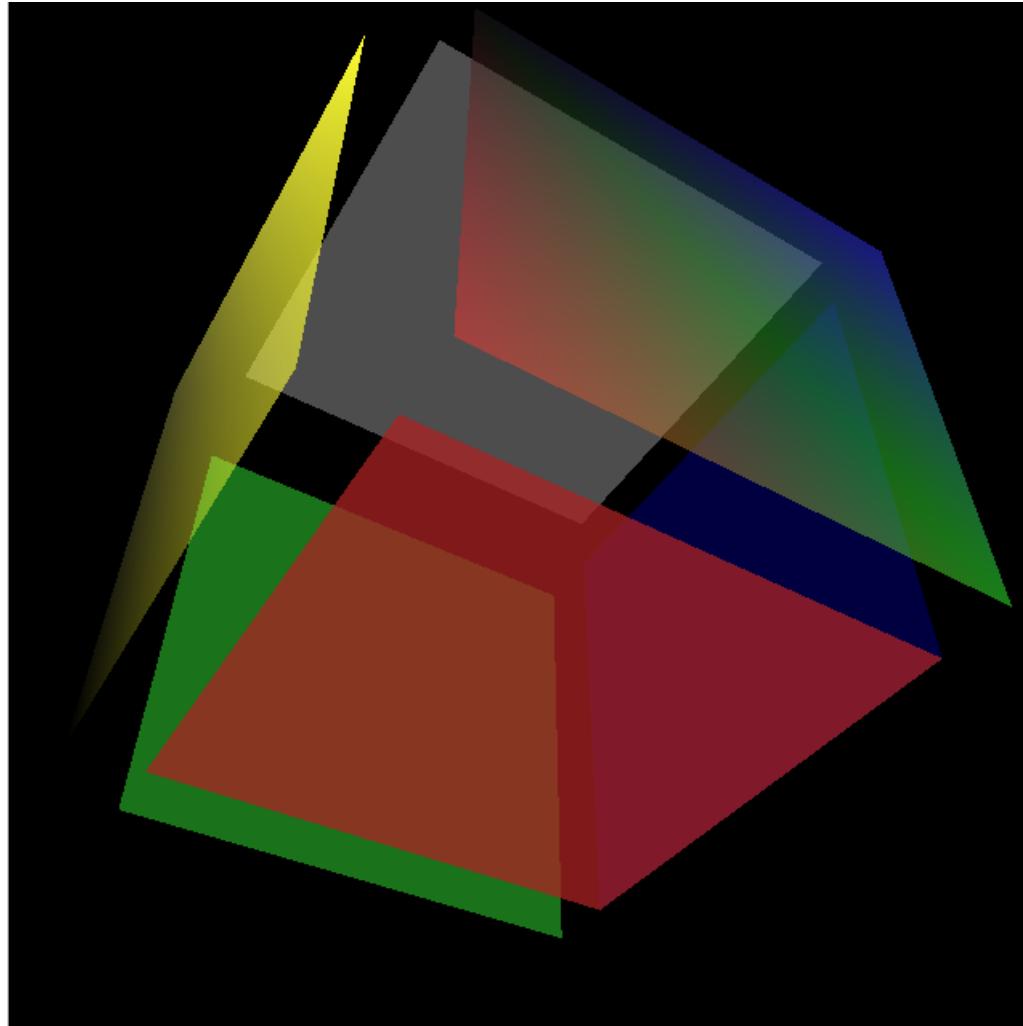


# Depth-compositing – VMD (UIUC)



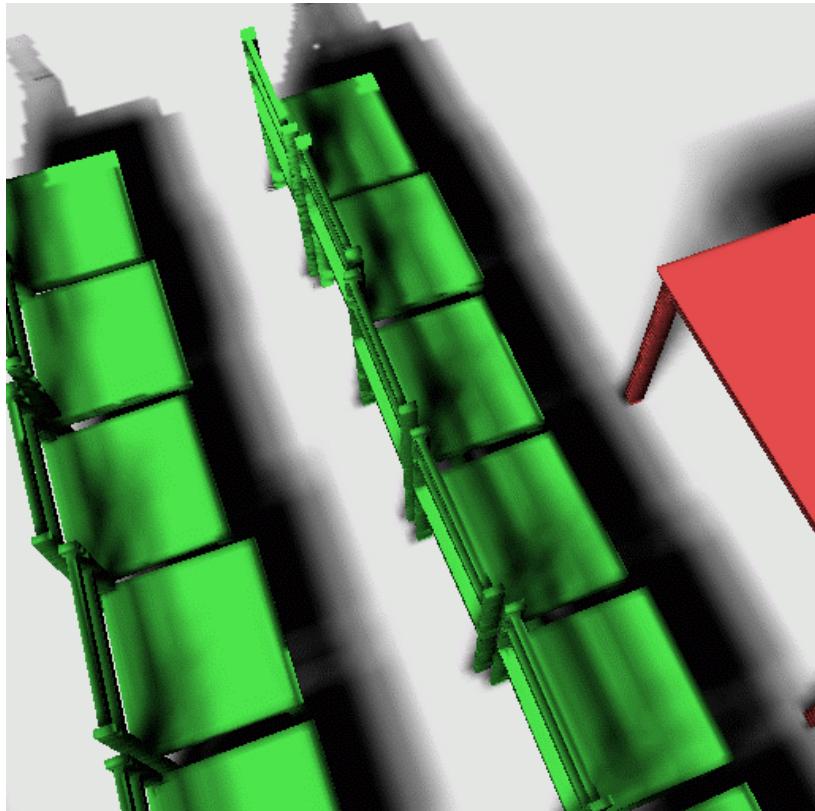
- Molecular visualization program for displaying, animating, and analyzing large biomolecular systems such as proteins
- Performance metric is frame rate of rotating view
- Simple parallelization, scaling depends on load balance

# Translucency (non-commutative)

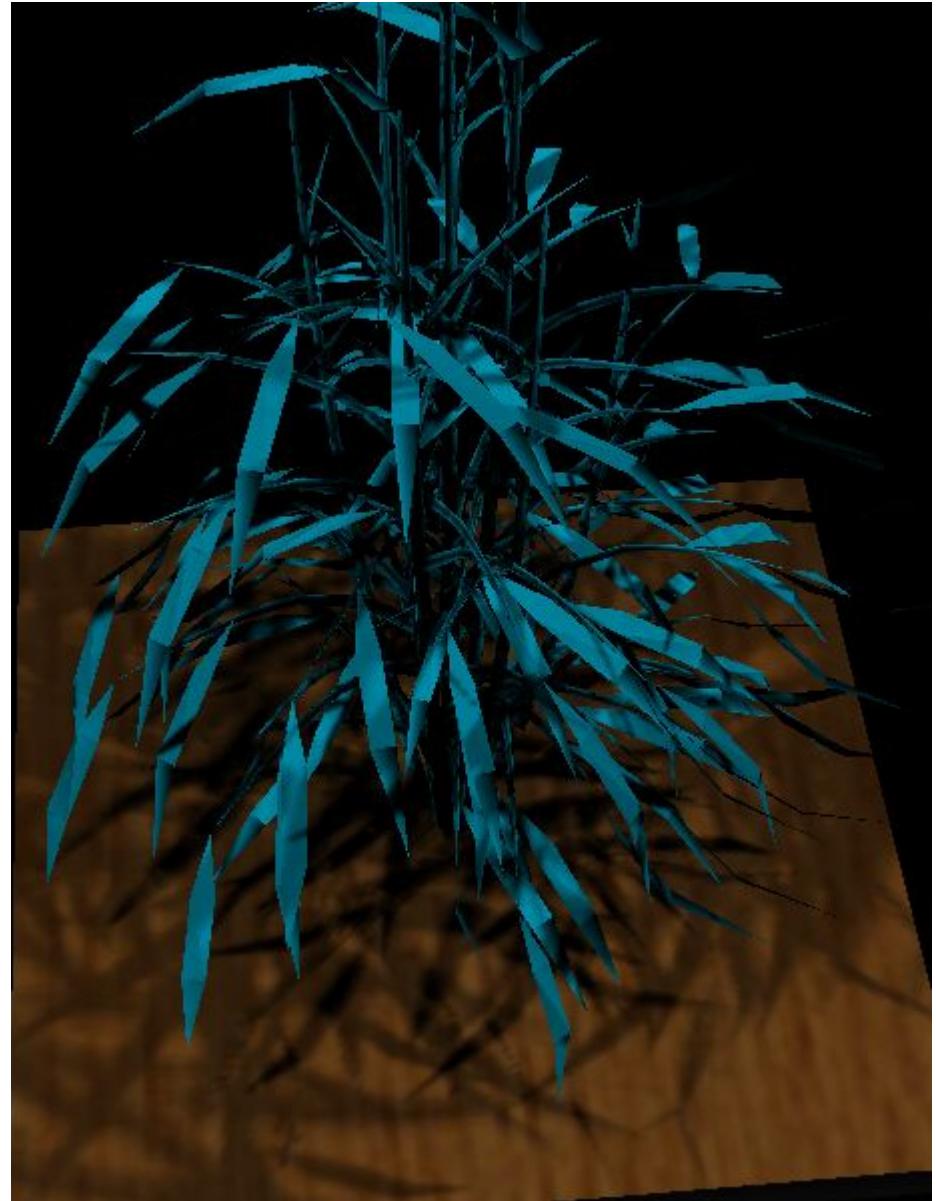


- Models light transmission
- Blending arithmetic is order-dependent (non-commutative)
- Concurrent evaluation requires associative formulation
- Equivalent to Porter-Duff translucency in Linux XFree86
- Solved by packet-switching

# Scalable lighting (soft shadows)



*Agrawala, Ramamoorthi, Moll & Heirich, SIGGRAPH 2000*



*Isard, Heirich and Shand, Eurographics PGV2002*



i n v e n t

